arXiv:1701.07011v1 [stat.ME] 24 Jan 2017

Comparable variable selection with Lasso

Lingfei Wang* and Tom Michoel

Department of Genetics & Genomics, The Roslin Institute
The University of Edinburgh, Easter Bush, Midlothian EH25 9RG, UK

January 25, 2017

Abstract

P-values are being computed for increasingly complicated statistics but lacking evaluations on their quality. Meanwhile, accurate p-values enable significance comparison across batches of hypothesis tests and consequently unified false discover rate (FDR) control. This article discusses two related questions in this setting. First, we propose statistical tests to evaluate the quality of p-value and the cross-batch comparability of any other statistic. Second, we propose a lasso based variable selection statistic, based on when the predictor variable first becomes active, and compute its p-value to achieve unified FDR control across multiple selections. In the end, we apply our tests on `covTest`, `selectiveInference`, and our statistic, based on real and null datasets for network inference in normal and high-dimensional settings. Results demonstrate higher p-value quality from our statistic and reveal p-value errors from others hidden before. We implement our statistic as `lassopv` in R.

Contents

1 Introduction **2**

2 Statistical tests for p-value quality **3**
2.1 Single-batch p-value tests . . . 3
2.2 Cross-batch p-value tests . . . 3
2.3 Practicalities for above tests . . . 4

3 Variable selection statistic based on the lasso regularization path **5**
3.1 Linear regression with lasso . . . 5
3.2 Variable selection with lasso . . . 5

4 Example applications and comparisons **7**
4.1 Data and methods to evaluate . . . 7
4.2 Evaluation results on reduced datasets . . . 8
4.3 Evaluation results on full datasets . . . 12

5 Discussion **19**

A Proofs **19**
A.1 Proposition 1 . . . 19
A.2 Proposition 2 . . . 20

*Corresponding author: Lingfei.Wang@roslin.ed.ac.uk

1 Introduction

Multiple hypothesis testing extensively uses p-values for false discovery rate (FDR) control, either in the form of frequentist p-value cutoff [1–3], or through the empirical estimation of local FDRs [4, 5]. In order to reach meaningful inference, both FDR control methods require accurate p-values as well as a null hypothesis that connects to the question. On one hand, most studies face relatively simple statistics, such as Pearson correlation whose p-value is easy to compute exactly and whose null hypothesis is easy to secure. On the other hand, there is a growing interest in more complicated statistics with realworld applications, yet in which those requirements are harder to satisfy, such as the sequential hypothesis testing where a null (or non-null) assertion for a hypothesis is simultaneously applied on its subsequent hypotheses.

For sure, an error-prone p-value would lead to less effective FDR control. For an experiment testing multiple hypotheses, it is most often that the p-value is a monotonic function of its statistic, therefore still leaving the same significance *ranking* for the hypotheses (or at most creating ties). However, when multiple hypotheses are tested by several different experiments, such as in meta-analysis [6, 7], different experimental conditions may introduce batch effects into their p-values. The same applies on other types of batches such as p-values from different variable selection tasks. Therefore, when comparing the significance levels of hypotheses from these different batches, batch-dependent errors in p-value estimation can invert the significance ranking despite individual monotonicity within each batch. As a result, errors in p-values would potentially lead to reduced statistical power in hypothesis testing when compared across multiple batches, thereby becoming more detrimental in addition to the less effective FDR control. In practice, there is also the need for accurate p-values that would allow a single significance measure to achieve an accurate, unified FDR control across multiple batches of hypothesis tests. It is therefore important to evaluate the quality of p-values especially for complicated statistics before drawing any inference. However, existing studies typically evaluated p-values in the simplest visual form — an observed-expected p-value plot for all hypotheses combined [8].

Section 2 covers the first intention of this paper, addressing this question with novel, quantitative and visual, statistical tests for p-value. We start from the uniform distribution of null p-values, and consider a large number of heterogeneous batches of hypotheses. In order to assess if estimated p-values can provide a unified FDR control across all batches, we propose to measure the linear correlation between the numbers of false positives and false cases, at any given p-value significance threshold. Derived statistics are also proposed subsequently, such as the R^2 goodness of fit, its curve as a function of the proportion of significant cases, and its area under the curve (AUC). We also apply the Kolmogorov-Smirnov (KS) test for each batch of hypotheses separately. These tests can also generalize and be applied on any generic statistic, including heavily biased p-value estimates. For these generic statistics, estimating the FDR level is impossible, but these tests can still evaluate whether a unified significance measure can be obtained across multiple batches without the knowledge of FDR.

The second question we face arises from variable selection, which can be regarded as testing the null hypothesis of each predictor variable not being selected. However, variable selection typically aims to select the least number of predictors but explain the most variance from target variable, therefore leading to the interdependency between the hypothesis testing outcomes of the predictors. This creates special difficulty in p-value calculation, especially for more complicated statistics which aim at better selection. Recent interests have arisen in obtaining accurate p-values in variable selection [8–10], or its FDR control [11]. Here we share the same interest, but aim at the consequence of accurate p-values — unified FDR control across multiple variable selections, every of which is an individual batch of hypothesis testing.

Lasso regularization of linear regression [12] has been extensively studied and applied in variable selection, especially when sparsity is desired from a large set of correlated variables. Various methods and statistics have been proposed, in which the simplest might just specify a regularization strength and select the active variables with nonzero coefficients. Recent methods and software packages have been devised specifically for p-value estimation or FDR control: `covTest` applies the covariance test statistic and its p-value to measure the local loss of explained variance when a variable is removed [9]; `selectiveInference` performs post-selection inference based on sequential hypothesis testing [8, 10]; `knockoff` introduces new variables to mimick the existing correlation

structure [13].

In Section 3, we propose a simple yet different statistic based on the naive assumption that variables that enter the lasso regularization path earlier tend to be more significant, measured by the regularization strength at which the variable becomes active for the first time. The advantage of this non-parametric statistic is its p-value that is easy and accurate to compute under analytical approximation, therefore enabling unified FDR control. We have released the method as the `lassopv` package in R, also available at https://github.com/lingfeiwang/lassopv.

In Section 4, the two objectives of this paper converge. We perform existing and novel evaluations on the p-values from existing and newly proposed lasso-based statistics. We consider the specific question of edge-pruning of a maximal Bayesian network originated from the inference of a gene regulation network, consisting of 3172 genes in 360 samples from the Geuvadis consortium [14]. Real and random null, normal and high-dimensional datasets are generated and considered based on the full dataset, in order to assess the quality of p-values from these statistics under a variety of conditions.

2 Statistical tests for p-value quality

2.1 Single-batch p-value tests

Consider n null p-values, namely $p_1, \ldots, p_n$, each of which is obtained from testing a generically different null hypothesis. In single-batch p-value tests, we do not use any batch information from these p-values. Therefore, they may either come from a single batch of hypothesis tests, or be combined from multiple batches. By definition, exact, null p-values should follow the independent identical standard uniform distribution $U(0,1)$, i.e.

$$p_i \sim \text{i.i.d } U(0,1), \text{ for } i = 1, \ldots, n. \tag{1}$$

Therefore, given n computed null p-values, we can empirically evaluate their quality by comparing them against the $U(0,1)$ distribution in a number of measures:

- Existing studies typically drew the observed-expected p-value plot as a visual clue (e.g. Figures 2 to 5 in [8]), in which the observed p-values are p_i, and the expected ones are their empirical cumulative distribution function (ECDF) values. Despite its wide adoption, the observed-expected p-value plot does not scale with sample size, produce quantitative measures, or clearly reveal the local deviations of the probability density function (PDF) from the standard uniform distribution. Therefore, we did not draw this plot in this article.
- **Histogram test**: Instead, we compared the histogram of computed null p-values with a standard uniform histogram of the same observation size. This is similar with the observed-expected p-value plot, but would allow a better visual inspection on potential PDF discrepancies. For overall comparison, we drew histogram on the combination of p-values from all batches of hypothesis tests.
- **The Kolmogorov-Smirnov test**: The KS test compares two (optionally empirical) cumulative distribution functions (CDFs) for their strongest discrepancy at every value, which provides an overview on the macroscopic disagreement between two (E)CDFs. We therefore applied the KS test on null p-values, separately for each batch of hypotheses, against the CDF of the standard uniform distribution. The KS statistic p-values were then shown in a Manhattan plot as a function of the number of hypotheses, and compared against the significance level 0.05 with Bonferroni correction to signify deviations from the standard uniform distribution.

2.2 Cross-batch p-value tests

In additional to the per-batch p-value tests, we also examined whether computed null p-values can obtain a unified FDR control, which is the motivation and prerequisite for many p-value applications. The principle is simple — if n computed p-values are ideal, i.e. each following the $U(0,1)$ distribution, then with any p-value significance threshold p_{thres}, the number of significant p-values would follow a binomial distribution $B(n, p_{\text{thres}})$. This approximates the normal distribution

$N(np_{\text{thres}}, np_{\text{thres}}(1-p_{\text{thres}}))$ when n is sufficiently large and, with multiple batches of hypotheses, we would see a linear relation between the numbers of p-values and significant p-values in each batch or, equivalently, between the numbers of false and false positive cases.

Based on this principle, we designed the cross-batch statistical tests for p-values as below:

- **Linear relation test**: With a given p-value significance threshold, we scatter-plotted the number of significant hypotheses of each batch as a function of the number of hypotheses of the same batch. A linear relation would indicate proportional false discoveries across all batches, and therefore guarantee a unified significance measure and FDR control, at least at this significance level. This was repeated at different significance thresholds to visually assess the linear relation.
- **Goodness of fit test**: Based on the linear relation test, we further performed linear regression to compute the goodness of linear fit in terms of R^2 on each linear relation. By computing R^2 at different significance thresholds, we drew R^2 as a function of the total proportion of significant hypotheses. A higher curve or a larger AUC thus indicates better unified significance measure and higher quality from the p-value. Since in most scenarios the accuracy of small p-values are more desirable to distinguish null and non-null cases, the R^2 values at smaller thresholds and correspondingly partial AUCs were computed.

  As extra information, we also included the maximum number of significant hypotheses across all batches. For accurate p-values and unified significance measures, the maximum number of significant hypotheses should grow linearly with the number of significant hypotheses. Therefore, it becomes a quality measure for p-value and also an indicator on the distribution of significant p-values. This will be utilized in high-dimensional variable selection in Section 4.3.

2.3 Practicalities for above tests

With any generic statistic, including heavily biased p-values, we can still obtain unified FDR control as long as they are unified significance measures, i.e. their (empirical) null distribution remains the same across all batches of hypothesis tests. On such occasions, we can ensure all batches reach (almost) the same FDR, although the FDR level remains unknown. Further knowing the null distribution of the statistic, such as the standard normal distribution for p-value, would enable the estimation of FDR level. Therefore, the above tests can also extend to generic statistics by removing the limitation of standard normal distribution. This disables the histogram test and modifies the KS test on the one-vs-rest basis — the ECDF of every batch is compared against the combined ECDF of all other batches instead — whilst the two cross-batch tests remain unchanged. These adjustments would allow comparison of unified FDR control power between p-values and generic statistics.

Besides, real data usually contain null as well as non-null hypotheses. The above tests only act on null p-values, which therefore we need to single out, for example with known groundtruths. On many occasions, however, groundtruths are unknown and that is the exact reason we perform hypothesis testing. Therefore, practical concerns must be taken to distinguish and remove the non-null hypotheses before proceeding with the statistical tests. The simplest option is to regard all hypotheses as null, if the proportion of non-null hypotheses is known to be very small. Another option is to remove all p-values deemed significant below a certain threshold, such as the conventional 0.05 without accounting for multiple testing, or even lower ones as long as the test statistics remain stable. The uniform distribution to test against should also be scaled accordingly, such as to $U(0.05, 1)$. For generic statistics, the significance threshold can be adjusted to the top/bottom 5%, for example, also if the proportion of non-null hypotheses is known to be small.

The cross-batch tests rely on the approximation from binomial to normal distributions. There is no doubt that more precise tests can be formulated without this approximation, but as we will show in the upcoming sections, the qualities of existing p-values already differed so much that the proposed tests were sufficient. For this reason, this article does not discuss any improvement on these tests.

3 Variable selection statistic based on the lasso regularization path

Variable selection is testing the null hypothesis of each variable not being selected, and every variable selection task can be regarded as a batch of hypotheses. Whenever we compare the significance of predictor variables from different variable selection tasks, the statistic (or its p-value) needs to be a unified signifcance measure or equivalently, achieve unified FDR control, which can be manifested in the above tests. In this section, we aim to develop a lasso-based statistic, and its p-value, to fulfill this purpose.

3.1 Linear regression with lasso

Given n observations of the target variable as $\mathbf{y} \in \mathbb{R}^n$ and k predictor variables as $\mathbf{X} \in \mathbb{R}^{n\times k}$, we assume the target and predictors are individually normalized to zero mean, i.e.

$$\sum_i y_i = 0, \tag{2}$$

$$\forall i = 1, \ldots, k, \qquad \sum_j x_{j,i} = 0. \tag{3}$$

We don't normalize the variance of $\mathbf{X}$ to allow for different assignments of relative regularization strengths, or that of $\mathbf{y}$ for the upcoming part of this paper.

The lasso regularization then solves the optimization problem

$$\hat{\beta}(\lambda) = \underset{\beta\in\mathbb{R}^k}{\operatorname{argmin}} \frac{1}{2n}||\mathbf{y} - \mathbf{X}\beta||_2^2 + \lambda||\beta||_1, \tag{4}$$

and predicts $\mathbf{y}$ with the estimator $\hat{\mathbf{y}}(\lambda)$ as

$$\hat{\mathbf{y}}(\lambda) = \mathbf{X}\hat{\beta}(\lambda). \tag{5}$$

To guarantee uniqueness of the regularization path, we assume the columns of $\mathbf{X}$ are in general position [15]. The parameter $\lambda > 0$ then decides the strength of regularization. Starting from the limit $\lambda \to \infty$, all $\hat{\beta}_i$'s are zero, leaving an empty *active set* of predictors. As we decrease λ gradually, predictors can enter or leave the active set, with the evolution of coefficients $\hat{\beta}(\lambda)$ forming the *regularization path*. The λ values at which a predictor is activated or deactivated are *knots*. Exact [16] and inexact [17] methods have been proposed to obtain the full regularization path with software packages readily available.

3.2 Variable selection with lasso

Consider k predictor variables $\mathbf{X} \equiv (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k)$ to be selected against the target variable $\mathbf{y}$, under the same setup as Section 3.1. Solving the lasso regression gives the regularization path as $\hat{\beta}(\lambda)$. Define

$$\lambda_i \equiv \sup\{\lambda : \hat{\beta}_i(\lambda) \neq 0\} \tag{6}$$

as the regularization strength at which variable $\mathbf{x}_i$ becomes active for the first time. Define the variance of $\mathbf{x}_i$ as

$$\sigma_i^2 \equiv \frac{1}{n}||\mathbf{x}_i||_2^2. \tag{7}$$

We define the random variable following the null hypothesis of predictor $\mathbf{x}_i$ as $\mathbf{x}_i^{(0)}$, a vector that started from *i.i.d* normal distribution with variance σ_i^2 for every element, but was then normalized to zero mean. We use superscript (0) for the null hypothesis of predictor $\mathbf{x}_i$, where all other predictors remain unchanged. The null hypothesis also leads to a new lasso regularization path $\hat{\beta}^{(0)}(\lambda)$. The test statistic is λ_i. The one sided p-value is then defined as the probability of test statistic under the null hypothesis being greater than the nominal value.

The p-value of variable $\mathbf{x}_i$ then becomes

$$P\left(\sup\{\lambda : \hat{\beta}_i^{(0)}(\lambda) \neq 0\} \geq \lambda_i\right) = P\left(\hat{\beta}_i^{(0)}(\lambda_i) \neq 0\right) + P\left(\sup\{\lambda : \hat{\beta}_i^{(0)}(\lambda) \neq 0\} \geq \lambda_i \wedge \hat{\beta}_i^{(0)}(\lambda_i) = 0\right). \tag{8}$$

The first term is the probability that $\mathbf{x}_i^{(0)}$ is active at $\lambda = \lambda_i$, whilst the second is the probability that $\mathbf{x}_i^{(0)}$ is active at some $\lambda > \lambda_i$ but inactive at $\lambda = \lambda_i$.

In lasso regression, predictors can enter and leave the active set multiple times. This leads to a nonzero value for the second term in Eq 8 in general. However, here we make and stick to the approximation of neglecting the second term, with justifications in Remark 1. It should be noted that by neglecting the second term in Eq 8, we would under-estimate the p-values which can lead to visible consequences.

Remark 1. *Justifications for neglecting the second term in Eq 8:*

- *Discovery-aimed variable selection focuses on selecting significant predictor variables, which are supposed to have small enough p-values. The second term in Eq 8 grows as more knots are gone through along the regularization path. Significant variables become active earlier, and therefore this approximation incurs much smaller errors on them than in average, especially when the number of genuine predictors is small.*
- *Although more accurate calculation of p-value without the above approximation can be possible, it is worthwhile to first investigate the performance of the statistic and its p-value estimation in this simple, approximated setting. With real and simulated data, this would help us decide if any more accurate approach is necessary.*

In order to disentangle $\mathbf{x}_i^{(0)}$ from the regularization path, we remove $\mathbf{x}_i^{(0)}$ from the set of predictors and construct a smaller problem as

$$\mathbf{X}^{(-i)} \equiv (\mathbf{x}_1, \ldots, \mathbf{x}_{i-1}, \mathbf{x}_{i+1}, \ldots, \mathbf{x}_k). \tag{9}$$

The new regularization path is defined as $\hat{\beta}^{(-i)}(\lambda)$. Define the variance of null variable $\mathbf{x}_i^{(0)}$ as

$$\left(\sigma_i^{(0)}\right)^2 \equiv \frac{1}{n}||\mathbf{x}_i^{(0)}||_2^2. \tag{10}$$

Define $\mathrm{Cor}(\cdot,\cdot)$ as the Pearson correlation function, and

$$\mathbf{y}_{\mathrm{res}}(\lambda) \equiv \mathbf{y} - \mathbf{X}^{(-i)}\hat{\beta}^{(-i)}(\lambda), \qquad \sigma_{\mathbf{y}_{\mathrm{res}}}^2(\lambda) \equiv \frac{1}{n}||\mathbf{y}_{\mathrm{res}}(\lambda)||_2^2 \tag{11}$$

as the residue of $\mathbf{y}$ and its variance for the smaller problem respectively.

Proposition 1.

$$\hat{\beta}_i^{(0)}(\lambda) \neq 0 \iff \lambda < \sigma_i^{(0)}\sigma_{\mathbf{y}_{\mathrm{res}}}(\lambda)\left|\mathrm{Cor}(\mathbf{x}_i^{(0)}, \mathbf{y}_{\mathrm{res}}(\lambda))\right|. \tag{12}$$

For proof, see appendix.

Proposition 2. *For $\mathbf{x} \in \mathbb{R}^n$ whose elements follow the i.i.d standard normal distribution, its Pearson correlation with another vector $\mathbf{y} \in \mathbb{R}^n$ follows the χ^2 distribution*

$$n\sigma_{\mathbf{x}}^2\mathrm{Cor}^2(\mathbf{x}, \mathbf{y}) \sim \chi^2(1), \tag{13}$$

where $\sigma_{\mathbf{x}}^2$ is the variance of $\mathbf{x}$.

For proof, see appendix.

Therefore we obtain the p-value for every predictor i as

$$\begin{aligned} & P\left(\sup\{\lambda : \hat{\beta}_i^{(0)}(\lambda) \neq 0\} \geq \lambda_i\right) \\ \approx\ & P\left(\hat{\beta}_i^{(0)}(\lambda_i) \neq 0\right) \\ =\ & P\left(\lambda_i < \sigma_i^{(0)}\sigma_{\mathbf{y}_{\text{res}}}(\lambda_i)\left|\text{Cor}\left(\mathbf{x}_i^{(0)}, \mathbf{y}_{\text{res}}(\lambda_i)\right)\right|\right) \\ =\ & P\left(n\left(\frac{\sigma_i^{(0)}}{\sigma_i}\right)^2 \text{Cor}^2\left(\frac{\mathbf{x}_i^{(0)}}{\sigma_i}, \mathbf{y}_{\text{res}}(\lambda_i)\right) > \frac{n\lambda_i^2}{\sigma_i^2\sigma_{\mathbf{y}_{\text{res}}}^2(\lambda_i)}\right) \\ =\ & Q\left(\frac{1}{2}, \frac{n\lambda_i^2}{2\sigma_i^2\sigma_{\mathbf{y}_{\text{res}}}^2(\lambda_i)}\right), \end{aligned} \tag{14}$$

where $Q(\cdot,\cdot)$ is the regularized gamma function, also one minus the CDF of χ^2 distribution.

Remark 2. *In univariate variable selection, the p-value reduces to that of Pearson correlation.*

4 Example applications and comparisons

In this section, we apply multiple lasso-based variable selection methods to estimate a sparse directed acyclic graph (DAG) when a natural ordering of nodes is given, and evaluate those methods with the statistical tests in Section 2. Estimating a sparse DAG or Bayesian network from observational data is an important problem when modelling the signalling pathways or gene regulatory networks of biological systems [18]. When a natural ordering of nodes is given, the problem reduces to a series of individual penalized regressions, where each node is regressed on its predecessors in the node ordering [11, 19]. Because the number of predictors is non-uniform across the different regressions, obtaining uniform variable selection p-values that allow to set a unified significance threshold across all regressions is a challenging problem. We will compare the variable selection statistic and its p-value, as we proposed in Section 3, against existing lasso-based variable selection methods, and demonstrate the non-uniform characteristics of the existing methods that are revealed by our new tests but hidden otherwise.

4.1 Data and methods to evaluate

The Geuvadis consortium collected gene expression levels and genotype variations from lymphoblastoid cell lines of 360 European individuals [14]. Although the true natural ordering of genes is, of course, unknown, 3172 of the 23,722 measured genes were found to possess at least one significant expression quantitative trait locus, i.e. a genetic variant in proximity of the gene and associated to variation of its expression across individuals. These genetic variants can be used to causally order pairs of genes using an instrumental variable approach [20]. We inferred the probability of every possible pairwise regulation with `Findr` [21], and constructed a greedy, maximal DAG by adding one directed regulation at a time in descending probability order. Edges that were to create a loop were discarded. Expression levels were preprocessed to follow a standard normal distribution by replacing every entry with the inverse function of its ECDF, separately for each gene. The resulting maximal DAG contained 3172 genes and 5,029,206 directed edges, and represented a natural ordering of genes based on the available genetic data.

Based on the inferred maximal DAG and expression levels of the 3172 genes, we then performed lasso regression of each gene on its predecessors in the DAG, and evaluated `lassopv` (our newly proposed method and package), `covTest`, and `selectiveInference` (at multiple λ's). The package `knockoff` was not attempted because it only outputs the set of selected variables at any FDR control level, but not the p-value for any of the variables, which limits and complicates the assessments we can perform.

The following 4 sub-datasets were used:

- The full, real dataset has the complete information of the expression levels and maximal DAG of the 3172 genes from 360 individuals.

- The full, null dataset was constructed based on the full, real dataset, in which every expression level is replaced with a random sample of independent standard normal distribution.
- The reduced, real dataset contains a subset of the genes and a maximal Bayesian subnetwork from the full, real dataset. Genes were ordered by network edge and the top 150 were selected for the 360 individuals, together with their maximal DAG.
- The reduced, null dataset started from the reduced, real dataset, and replaced every expression level with a random sample of independent standard normal distribution.

The reduced and full datasets were aimed to assess method performances in normal and high-dimensional settings. Comparisons of method performances for null and real datasets also provide better completeness.

For every dataset, we computed the p-values for edge-pruning as follows. We regarded each gene as the target variable and all its potential regulator genes specified in the maximal DAG as predictor variables, thereby carrying out one variable selection per target gene. P-values of all potential regulations could then be obtained with the variable selection method/software. Using tests in Section 2, we then evaluated the quality of p-values from different variable selection methods and different datasets. The actual pruning step requires specifying a p-value threshold, such as through empirical FDR control [4, 5], which is off the topic of this paper and would not be discussed here.

4.2 Evaluation results on reduced datasets

In this section, we present our evaluation results for the variable selection methods on the reduced, low-dimensional datasets. All statistical tests in Section 2 were performed but results of the linear relation test will be omitted. This is because the limited number of predictor variables could not provide an obvious visual comparison. Instead, we will only present its downstream analysis on goodness of fit.

- **Histogram test**: For every method, we drew its computed p-value histograms on the reduced, null dataset in Figure 1 and on the reduced, real dataset in Figure 2. Our method `lassopv` produced uniformly distributed p-values on null data in Figure 1(a), on top of which an abundance close to zero could be observed for real data in Figure 2(a). On the other hand, `covTest` was strongly over-conservative in p-value estimation on both datasets. The performance of `selectiveInference` depended on the choice of regularization strength: strong regularization lead to over-abundance at small and large p-values on both datasets, whilst weak regularization provided uniformly distributed p-values on both datasets. This was in agreement with Figure 3 and 5 in [8], which were for logistic regression and Cox models though. The indifference of p-value histograms from `selectiveInference` on null and real datasets raised the question on `selectiveInference`'s ability to identify true predictors, as well as on the correct choice of λ.
- **The Kolmogorov-Smirnov test**: Results of the KS test on the reduced null dataset, for `lassopv` and `selectiveInference` with different λ's, as well as one-vs-rest KS test for `covTest` due to its biased p-values, are shown in Figure 3.

  Figure 3(a) exhibited ideal performances from `lassopv` under the KS test, suggesting null p-values consistently following $U(0, 1)$ distribution. However, `selectiveInference` did not produce p-values following $U(0, 1)$ distribution, regardless of the choice of λ, or after retreating to one-vs-rest KS test (not shown). Its KS statistics were inflated in general, with some KS p-values as small as 10^{-10}. The `covTest` also failed the one-vs-rest KS test in Figure 3(b), particularly at small variable counts.

  The reduced real dataset were then applied in the KS test evaluation. We attempted to recover the (sub)set of null variables by removing the p-values smaller than 0.01. Similarly for `covTest`, we removed the bottom 1% of its p-values and performed one-vs-rest KS test. Higher thresholds were attempted but lead to no significant difference in the outcomes. The evaluation results on the real dataset in Figure 4 were consistent with those on the null dataset.
- **Goodness of fit test**: We plotted R^2 and the maximum number of significant predictors as functions of the total proportion of significant hypotheses at different significance thresholds,

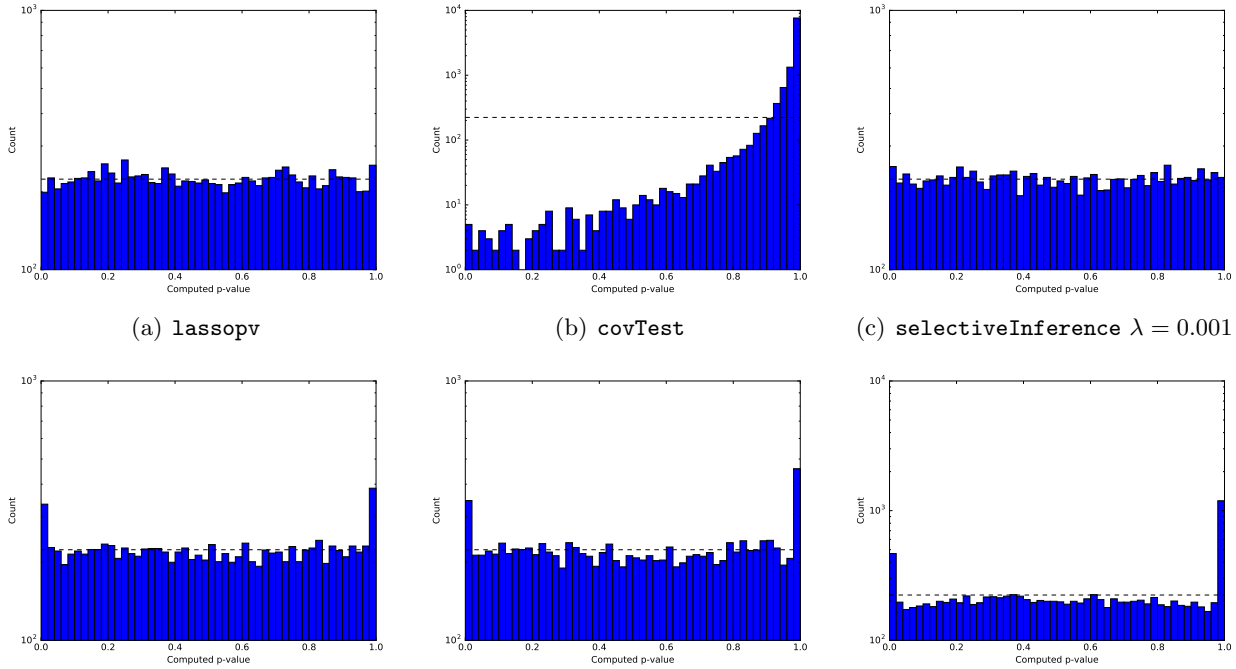

(a) `lassopv` (b) `covTest` (c) `selectiveInference` $\lambda = 0.001$

(d) `selectiveInference` $\lambda = 0.005$ (e) `selectiveInference` $\lambda = 0.02$ (f) `selectiveInference` $\lambda = 0.1$

Figure 1: The histogram of computed p-values on the reduced null dataset from `lassopv`, `covTest`, and `selectiveInference` at different λ values. The dashed line stands for the perfectly uniform histogram of the same size.

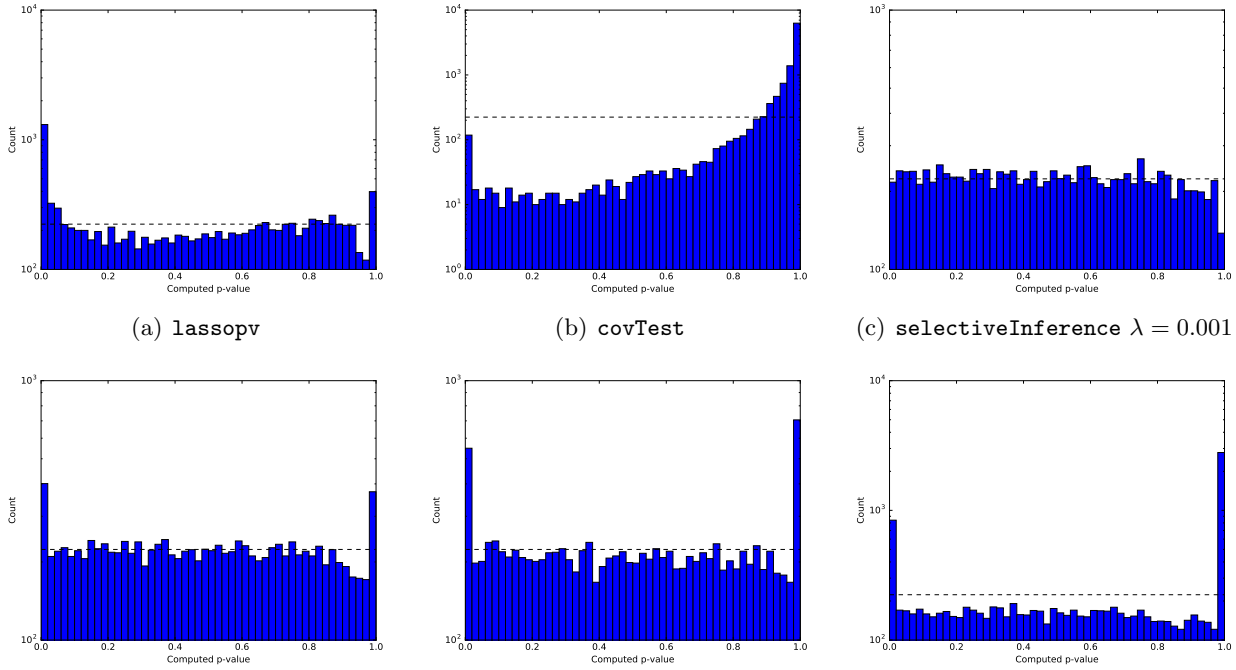

(a) `lassopv` (b) `covTest` (c) `selectiveInference` $\lambda = 0.001$

(d) `selectiveInference` $\lambda = 0.005$ (e) `selectiveInference` $\lambda = 0.02$ (f) `selectiveInference` $\lambda = 0.1$

Figure 2: The histograms of computed p-values on the reduced real dataset from `lassopv`, `covTest`, and `selectiveInference` at different λ values. The dashed line stands for the perfectly uniform histogram of the same size.

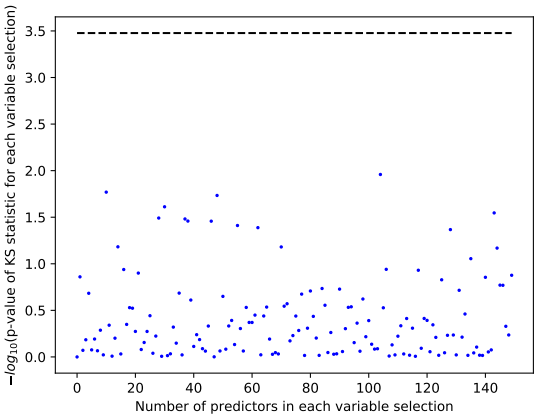

(a) `lassopv`

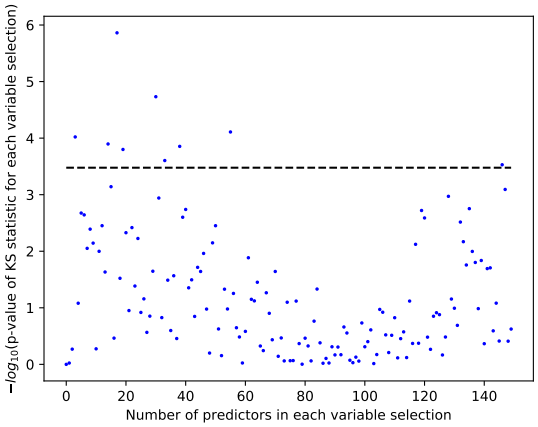

(b) `covTest` (one-vs-rest)

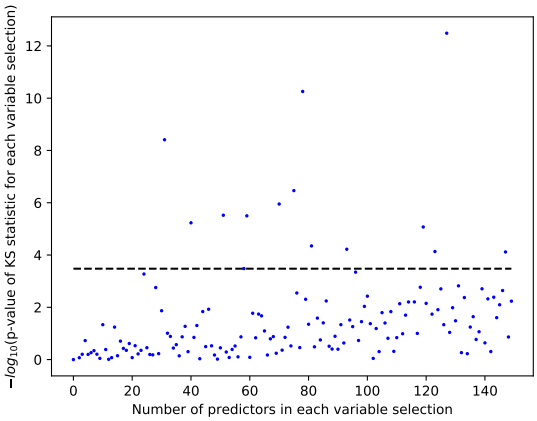

(c) `selectiveInference` $\lambda = 0.001$

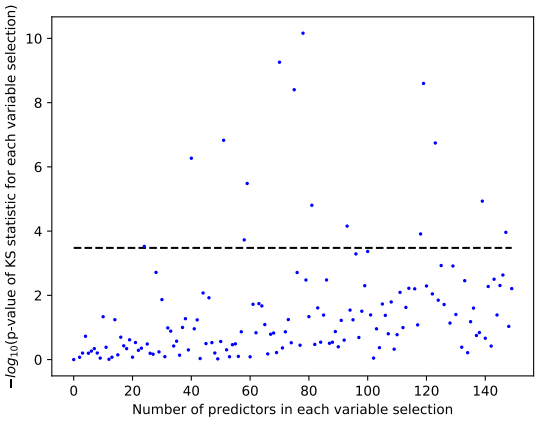

(d) `selectiveInference` $\lambda = 0.005$

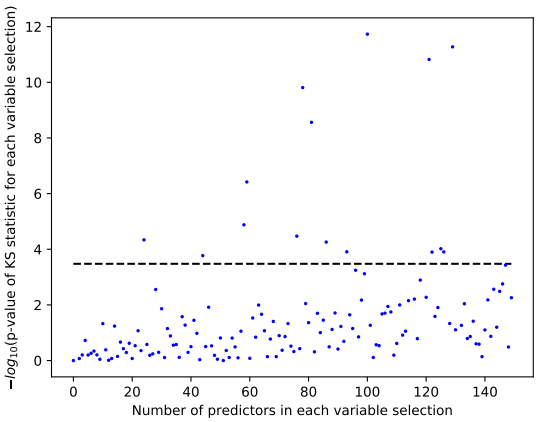

(e) `selectiveInference` $\lambda = 0.02$

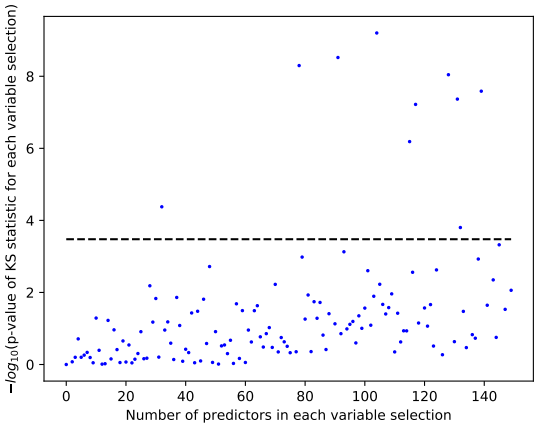

(f) `selectiveInference` $\lambda = 0.1$

Figure 3: For every of the 149 variable selection tasks of the reduced null dataset, we performed the KS test on its p-values against the $U(0, 1)$ distribution. We drew a Manhattan plot for the KS statistic p-value, together with the dashed line as Bonferroni corrected significance level 0.05 for deviation from the designated distribution. This is done separately for `lassopv`, `covTest`, and `selectiveInference` at different λ values. Due to `covTest`'s apparent deviation from uniform distribution in Figure 1(b), here we performed KS test on one-vs-rest basis instead.

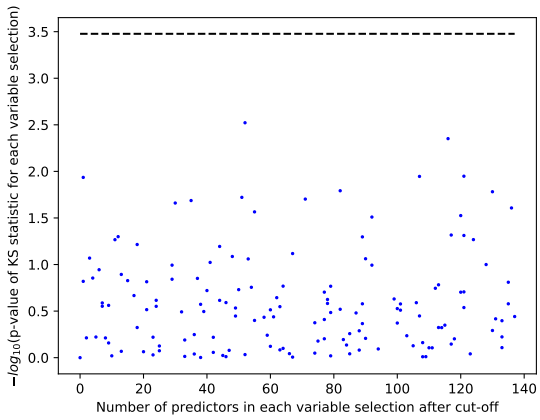
$-log_{10}$(p-value of KS statistic for each variable selection)
Number of predictors in each variable selection after cut-off

(a) `lassopv`

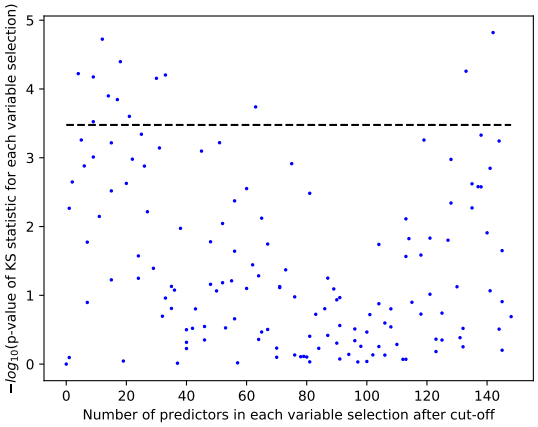
$-log_{10}$(p-value of KS statistic for each variable selection)
Number of predictors in each variable selection after cut-off

(b) `covTest` (one-vs-rest)

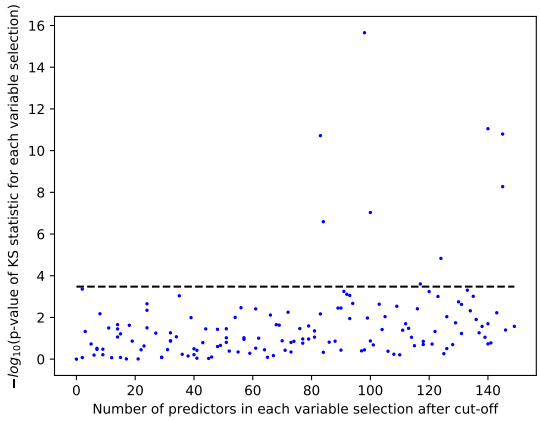
$-log_{10}$(p-value of KS statistic for each variable selection)
Number of predictors in each variable selection after cut-off

(c) `selectiveInference` $\lambda = 0.001$

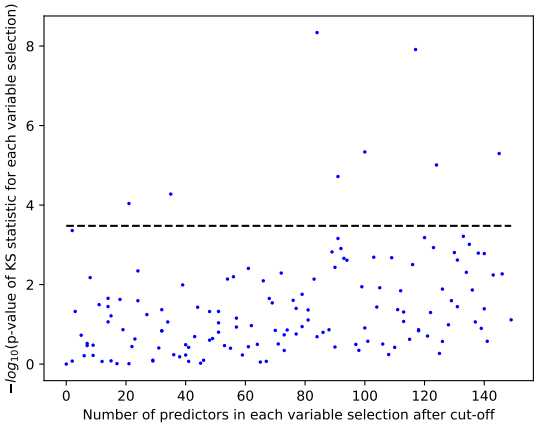
$-log_{10}$(p-value of KS statistic for each variable selection)
Number of predictors in each variable selection after cut-off

(d) `selectiveInference` $\lambda = 0.005$

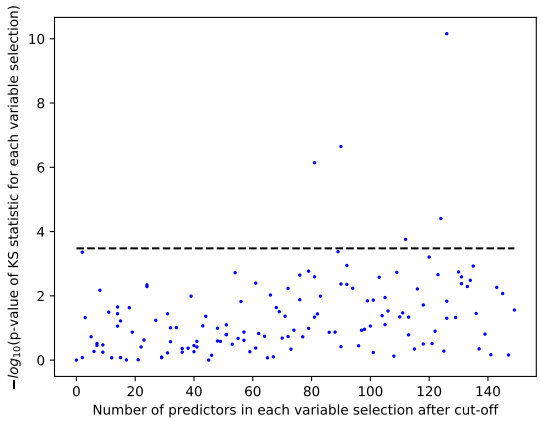
$-log_{10}$(p-value of KS statistic for each variable selection)
Number of predictors in each variable selection after cut-off

(e) `selectiveInference` $\lambda = 0.02$

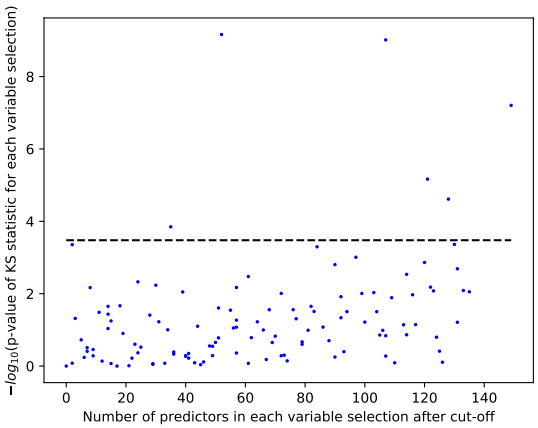
$-log_{10}$(p-value of KS statistic for each variable selection)
Number of predictors in each variable selection after cut-off

(f) `selectiveInference` $\lambda = 0.1$

Figure 4: The same KS test was performed on the reduced real dataset as in Figure 3. Without knowledge of groundtruth, we attempted to recover (a subset of) the null variables by excluding all variables with p-value smaller than 0.01, and compared the remaining p-values with $U(0.01, 1)$. For `covTest`, the p-values in the bottom 1% were excluded, after which one-vs-rest KS test was performed.

Table 1: Normalized partial and full AUCs are shown for the R^2 curves in Figure 5 and Figure 6. Partial AUCs were computed within the given x coordinate bounds (column *Bound*), and normalized to achieve unit AUC for constant function $R^2 = 1$. Best performers are shown in bold.

Bound	`lassopv`	`covTest`	`selectiveInference` λ =0.001	0.005	0.02	0.1
Null dataset						
$[0, 0.01]$	**0.136**	0.005	0.005	0.008	0.009	0.016
$[0, 0.05]$	**0.36**	0.04	0.02	0.03	0.04	0.07
$[0, 0.2]$	**0.68**	0.26	0.17	0.17	0.18	0.20
$[0, 1]$	**0.91**	0.77	0.74	0.72	0.72	0.60
Real dataset						
$[0, 0.01]$	0.047	0.004	**0.065**	0.008	0.012	0.015
$[0, 0.05]$	**0.18**	0.02	0.04	0.03	0.05	0.07
$[0, 0.2]$	**0.44**	0.20	0.18	0.16	0.19	0.20
$[0, 1]$	**0.84**	0.74	0.72	0.71	0.67	0.41

for the reduced null dataset in Figure 5 and reduced real dataset in Figure 6. The total proportion of significant hypotheses is equal to the FPR in the null dataset, and approximates the FPR in the real dataset due to the sparsity nature of gene regulation networks. Full and partial AUCs for R^2 were also computed in Table 1.

We observed consistent stronger linear relations from `lassopv` than `selectiveInference` and `covTest`, in terms of full and partial AUCs, whilst `selectiveInference` tended to assign small p-values to variables in several selection tasks in null and real datasets.

On real data, existence of the true and the correlated predictors are the partial causes of the declined performances. Note the exception that `selectiveInference` with $\lambda = 0.001$ outperformed other methods in AUC for $x \in [0, 0.01]$ in Table 1 and its R^2 spike close to zero in Figure 6(c). This should not be interpreted as good performance but, on the contrary, is due to its failure in identifying genuine predictors, as already suggested in Figure 2(c).

To summarize this section, `lassopv` performed much better than `covTest` and `selectiveInference` in the three evaluations on the null and real datasets of low dimension. `lassopv` performed well in all three tests. `covTest` failed to yield uniformly distributed null p-values. `selectiveInference` produced visually uniform overall histograms, but failed the Kolmogorov-Smirnov test. Neither `covTest` or `selectiveInference` achieved satisfying R^2 necessary for the unified FDR control across multiple variable selection tasks. In addition, `selectiveInference` suffered doubts in the choice of regularization strength and its statistical power to identify genuine predictor variables.

4.3 Evaluation results on full datasets

In this section, we present the results and remarks of the same evaluations but on the full high-dimensional (null and real) datasets for `lassopv` and `covTest`. The lasso regularization path was obtained with default number of knots in `lars` [16]. The package `selectiveInference` was attempted but failed critically, requesting 'higher precision' of the regularization path although the lasso regression was completed by `lars` under the same configuration. Hence we came to the understanding that the package `selectiveInference` cannot handle high-dimensional datasets of similar sizes.

- **Histogram test**: In high-dimensional lasso regression computation, the number of variables that ever become active is limited by the numbers of observations/samples and knots. Variables that never enter the active set should always have p-values of 1, which can be the majority depending on the dataset. Therefore, we specifically removed p-values of 1 before producing the histograms for the null and real datasets in Figure 7.

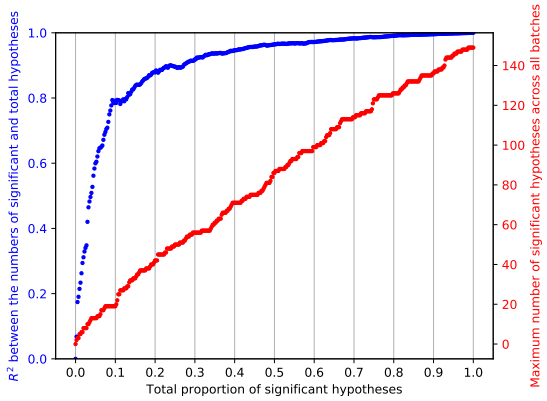
R^2 between the numbers of significant and total hypotheses
Maximum number of significant hypotheses across all batches
Total proportion of significant hypotheses

(a) `lassopv`

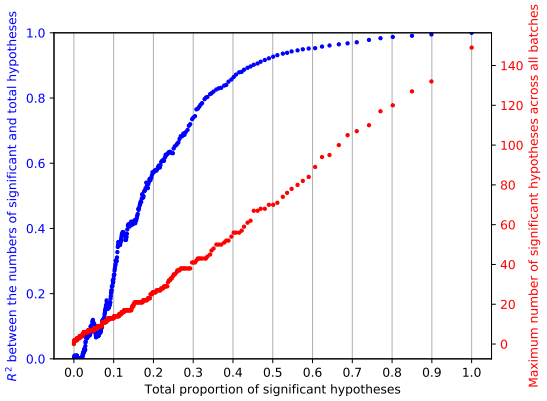
R^2 between the numbers of significant and total hypotheses
Maximum number of significant hypotheses across all batches
Total proportion of significant hypotheses

(b) `covTest`

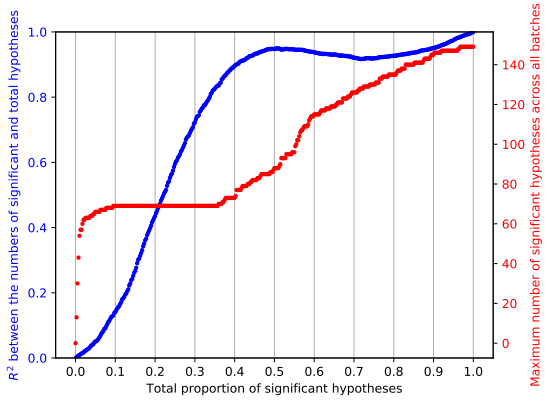
R^2 between the numbers of significant and total hypotheses
Maximum number of significant hypotheses across all batches
Total proportion of significant hypotheses

(c) `selectiveInference` $\lambda = 0.001$

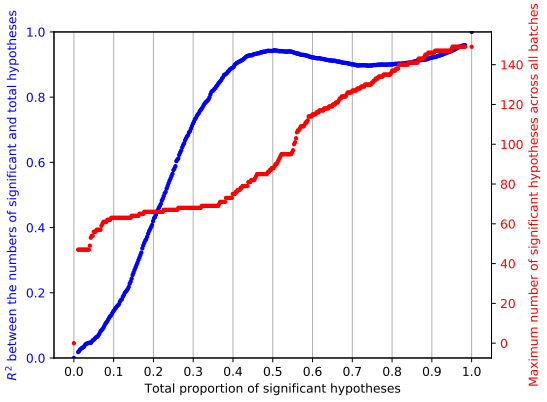
R^2 between the numbers of significant and total hypotheses
Maximum number of significant hypotheses across all batches
Total proportion of significant hypotheses

(d) `selectiveInference` $\lambda = 0.005$

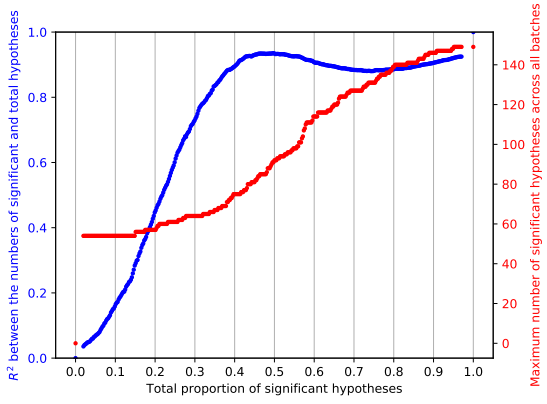
R^2 between the numbers of significant and total hypotheses
Maximum number of significant hypotheses across all batches
Total proportion of significant hypotheses

(e) `selectiveInference` $\lambda = 0.02$

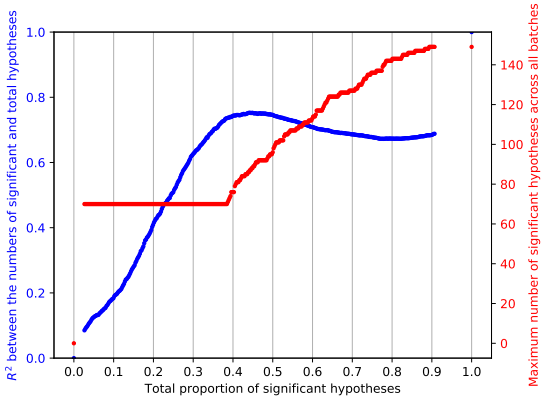
R^2 between the numbers of significant and total hypotheses
Maximum number of significant hypotheses across all batches
Total proportion of significant hypotheses

(f) `selectiveInference` $\lambda = 0.1$

Figure 5: The R^2 goodness of fit is shown in blue as a function of the proportion of significant hypotheses for the reduced null dataset. The maximum number of significant variables across all selection tasks is also drawn in red. Every dot corresponds to a significance threshold taken. Larger R^2 indicates better linear relation and therefore better unified FDR control.

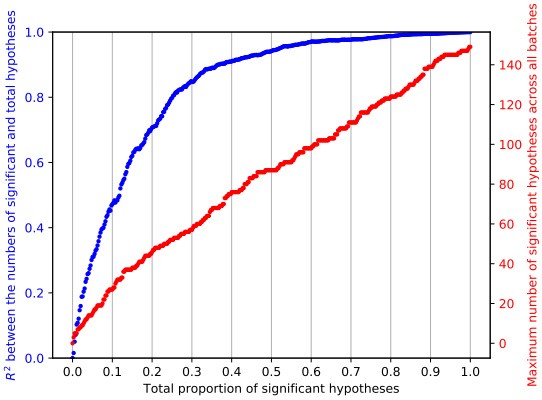
R^2 between the numbers of significant and total hypotheses
Maximum number of significant hypotheses across all batches
Total proportion of significant hypotheses

(a) `lassopv`

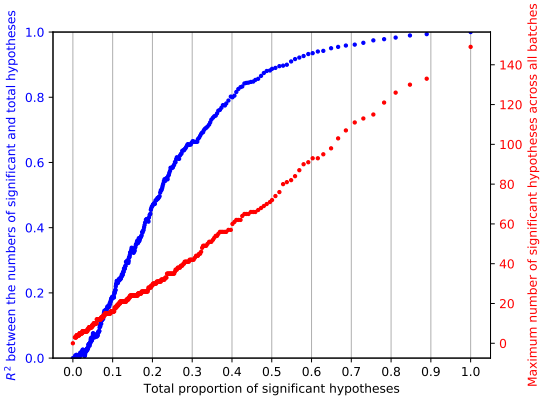
R^2 between the numbers of significant and total hypotheses
Maximum number of significant hypotheses across all batches
Total proportion of significant hypotheses

(b) `covTest`

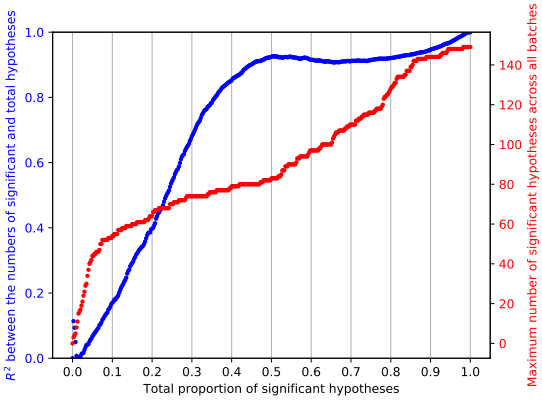
R^2 between the numbers of significant and total hypotheses
Maximum number of significant hypotheses across all batches
Total proportion of significant hypotheses

(c) `selectiveInference` $\lambda = 0.001$

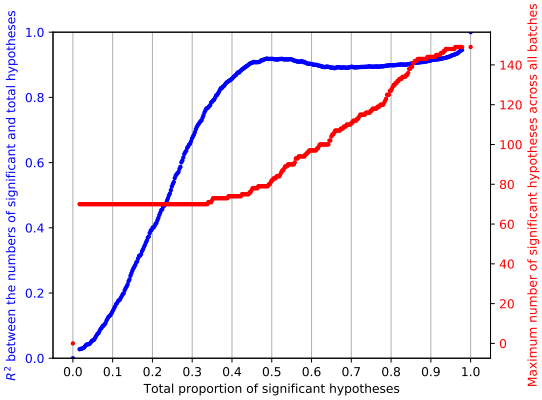
R^2 between the numbers of significant and total hypotheses
Maximum number of significant hypotheses across all batches
Total proportion of significant hypotheses

(d) `selectiveInference` $\lambda = 0.005$

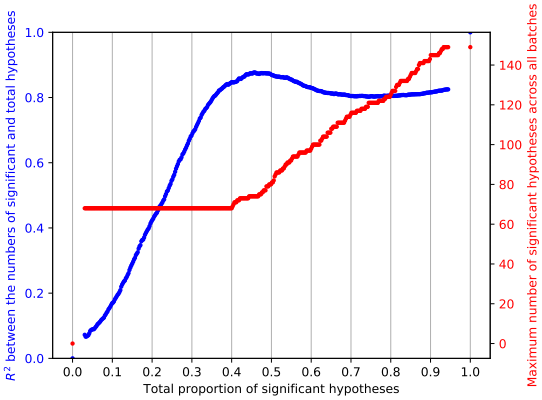
R^2 between the numbers of significant and total hypotheses
Maximum number of significant hypotheses across all batches
Total proportion of significant hypotheses

(e) `selectiveInference` $\lambda = 0.02$

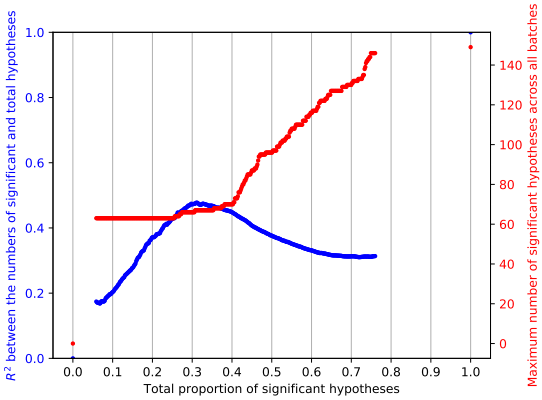
R^2 between the numbers of significant and total hypotheses
Maximum number of significant hypotheses across all batches
Total proportion of significant hypotheses

(f) `selectiveInference` $\lambda = 0.1$

Figure 6: The R^2 goodness of fit is shown in blue as a function of the proportion of significant hypotheses for the reduced real dataset. The maximum numbers of selected variables in one selection task were also drawn in red. Every dot corresponds to a threshold taken.

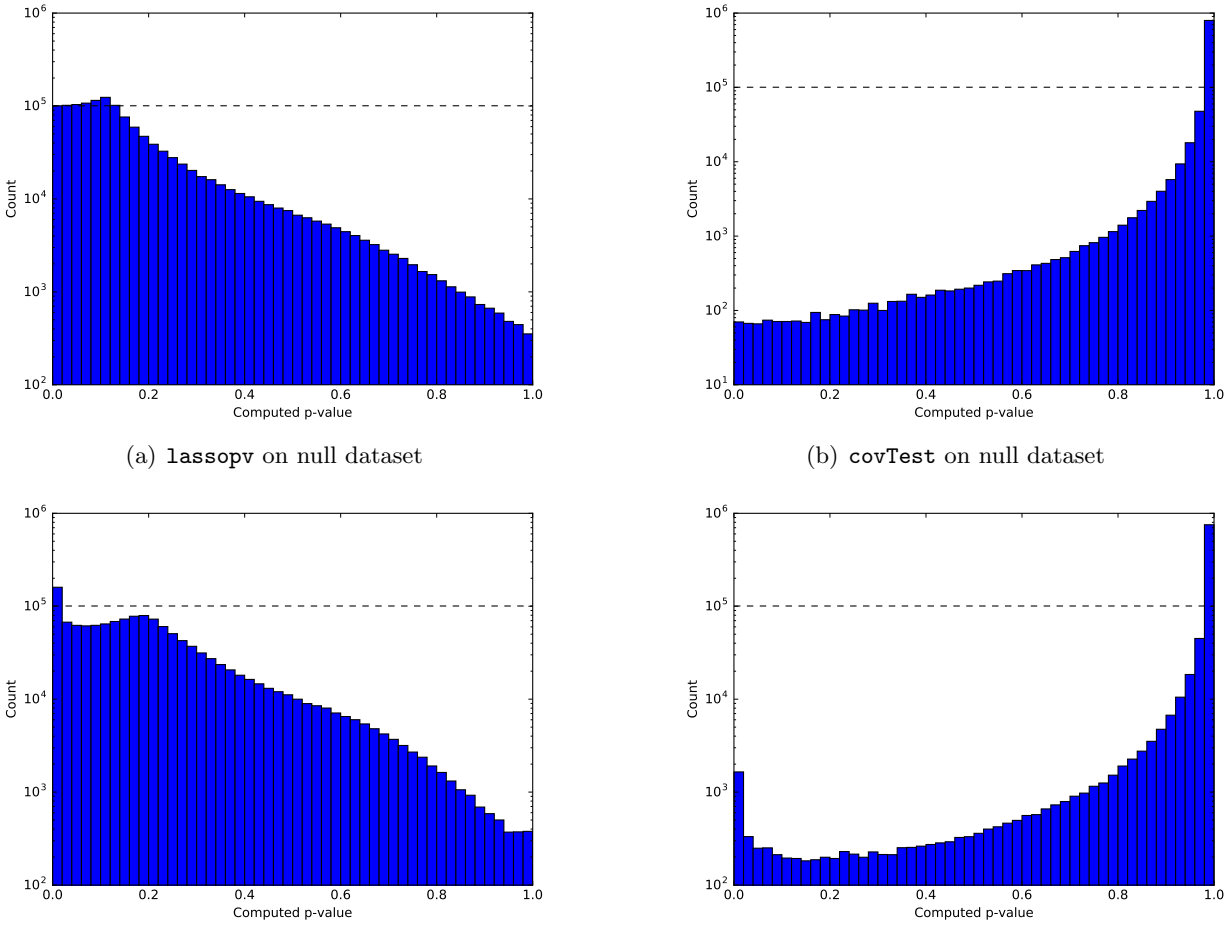

(a) `lassopv` on null dataset

(b) `covTest` on null dataset

(c) `lassopv` on real dataset

(d) `covTest` on real dataset

Figure 7: The p-value histograms from `lassopv` and `covTest` are plotted on the full null and real datasets, after the removal of p-values of 1. The dashed line stands for the perfectly uniform histogram of the same size before removing p-values of 1.

In Figure 7(a), we observed a significant deviation from uniform distribution by the p-values of `lassopv`. We believe there were two reasons:

- The dilution of large p-values was mostly due to variables that did not enter the computed regularization path within the given knots, and were therefore discarded. This is a problem solely of the computation of regularization path, shared by all variable selection statistics using lasso, and cannot be solved by `lassopv`.
- The over-abundance of p-values between 0.1 and 0.2 was partly due to the same high-dimensional effect, and partly to the approximation in Section 3.2 which under-estimated p-values as we proceeded along the regularization path. This was confirmed by a simulation-based p-value computation, which avoided the approximation and shrank the peak height by half for the real dataset (not shown).

In spite of that, properties of the two reasons ensure histogram distortions be very minor at small p-values, as confirmed in Figure 7, thus allowing for accurate estimation of p-values and FDRs for high-dimensional, yet sparse variable selection.

On the other hand, `covTest` remained over-estimating p-values for high-dimensional datasets.

- **The Kolmogorov-Smirnov test**: In order to minimize high-dimensional effects and approximation errors, and to target the most interested small p-values, we only performed the KS test on the most significant 5% p-values of the null dataset. The bound was modified to

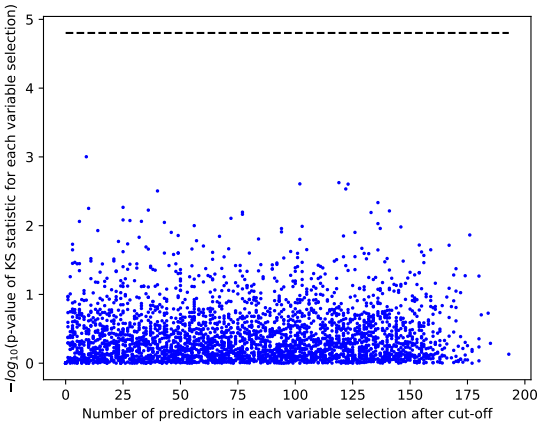
$-log_{10}$(p-value of KS statistic for each variable selection)
Number of predictors in each variable selection after cut-off

(a) `lassopv` on null dataset

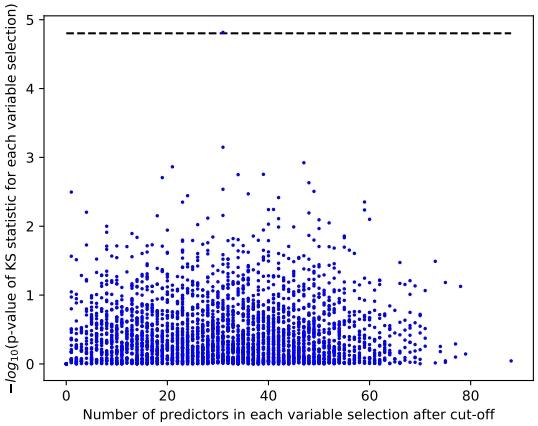
$-log_{10}$(p-value of KS statistic for each variable selection)
Number of predictors in each variable selection after cut-off

(b) `lassopv` on real dataset

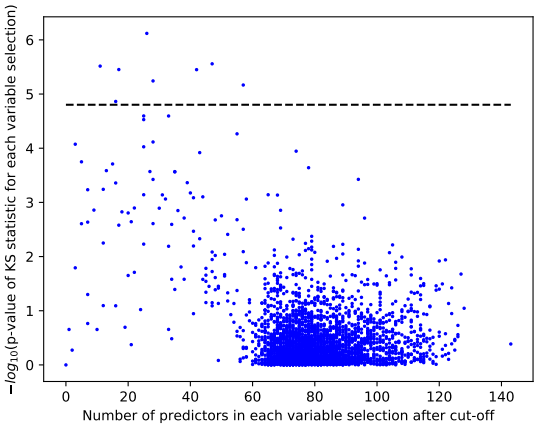
$-log_{10}$(p-value of KS statistic for each variable selection)
Number of predictors in each variable selection after cut-off

(c) `covTest` on null dataset (one-vs-rest)

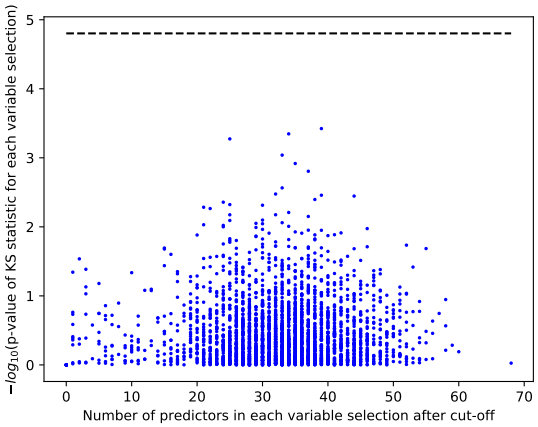
$-log_{10}$(p-value of KS statistic for each variable selection)
Number of predictors in each variable selection after cut-off

(d) `covTest` on real dataset (one-vs-rest)

Figure 8: The KS statistic p-values are shown for `lassopv` and `covTest` on high-dimensional null and real datasets. Different p-value bounds were applied to focus on small p-values whilst avoiding non-null cases — the most significant 5% for the null dataset and 3% to 5% for the real. In addition, one-vs-rest KS tests were performed for `covTest`.

between 3% and 5% for the real dataset to remove non-null cases. For `covTest`, one-vs-rest KS tests were performed to account for its non-uniform p-value distribution. As shown in Figure 8, KS test results are in agreement with the reduced, low-dimensional setting.

- **Linear relation test**: The linear relation scatter-plots are shown in Figure 9 for the null dataset. The linear relation between the numbers of predictors and significant predictors, as observed in Figure 9(a) and Figure 9(b), confirmed the unified significance measure from `lassopv` across all variable selection tasks. At high p-value thresholds, the number of significant predictors saturated as expected, leading to the plateau in Figure 9(c). On the other hand, `covTest` failed to provide a linear relation or consequently a unified significance measure. On the real, correlated dataset, the linear relation and R^2 were both less significant from `lassopv` and `covTest`, due to the introduction of non-null predictors, but `lassopv` still obtained strong linear relations, with details in the goodness of fit test.
- **Goodness of fit test**: As shown in Figure 10 and Table 2, `lassopv` overperformed `covTest` on both datasets with stronger linear relations that indicated better unified significance measure. The goodness of fit of `lassopv` decreased on real data as expected, due to the presence of genuine predictors.

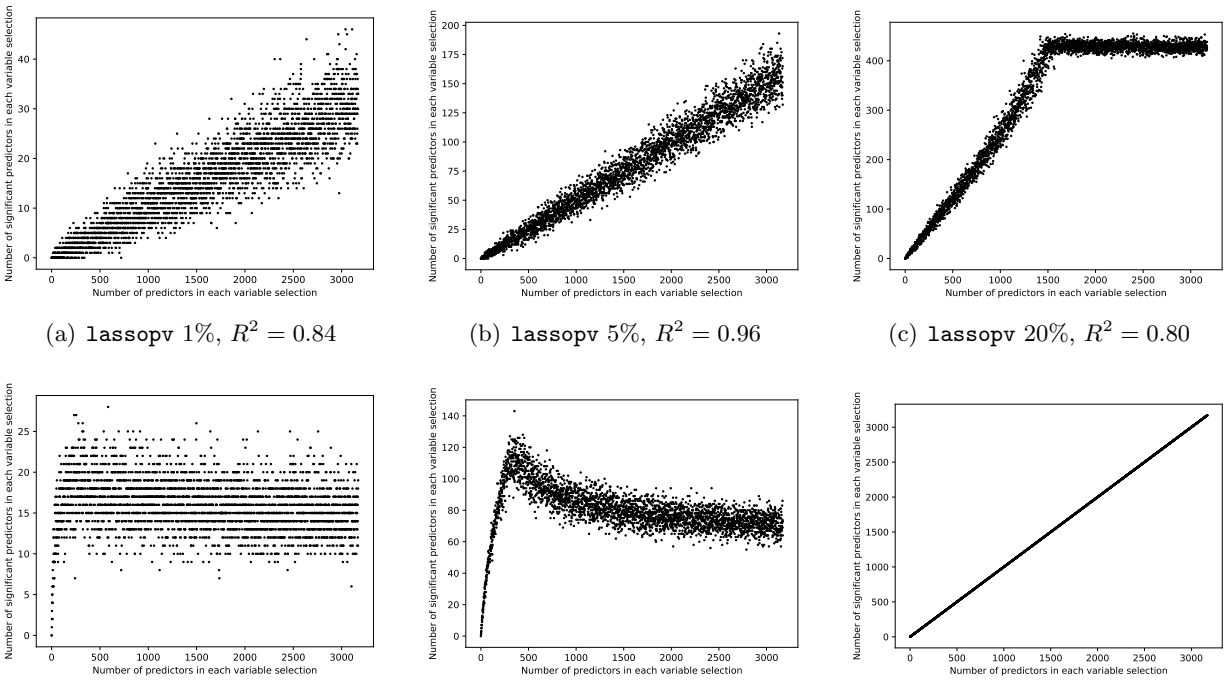

(a) `lassopv` 1%, $R^2 = 0.84$ (b) `lassopv` 5%, $R^2 = 0.96$ (c) `lassopv` 20%, $R^2 = 0.80$

(d) `covTest` 1%, $R^2 = 0.04$ (e) `covTest` 5%, $R^2 = 0.12$ (f) `covTest` 20%, $R^2 = 1$ (no variable selection)

Figure 9: The numbers of predictors and significant predictors are plotted at different p-value significance thresholds for `lassopv` and `covTest` on the full null dataset. The significance thresholds were chosen as the first 100-, 20- and 5-quantiles across all selection tasks, or equivalently the most significant 1%, 5%, and 20% p-values. Figure 9(f) has a straight line because all variables were regarded significant despite the 20% cutoff (as confirmed in Figure 10(c)). The goodness of fit in terms of R^2 is also shown.

Table 2: Normalized partial AUCs are shown for the R^2 curves in Figure 10. Best performers are shown in bold.

	Null dataset		**Real dataset**	
Bound	`lassopv`	`covTest`	`lassopv`	`covTest`
$[0, 0.001]$	**0.192**	0.0001	**0.019**	0.008
$[0, 0.01]$	**0.65**	0.01	**0.05**	0.03
$[0, 0.1]$	**0.93**	0.08	**0.59**	0.04

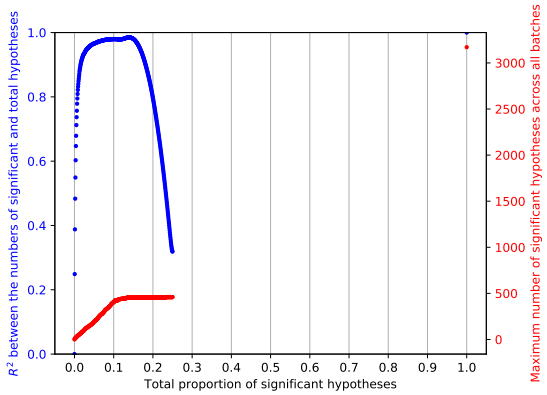
R^2 between the numbers of significant and total hypotheses
Maximum number of significant hypotheses across all batches
Total proportion of significant hypotheses

(a) `lassopv` on null dataset

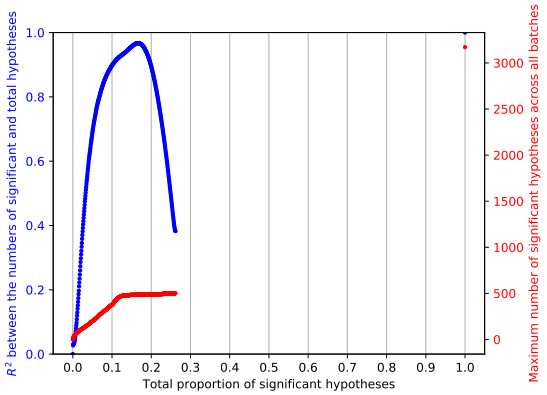
R^2 between the numbers of significant and total hypotheses
Maximum number of significant hypotheses across all batches
Total proportion of significant hypotheses

(b) `lassopv` on real dataset

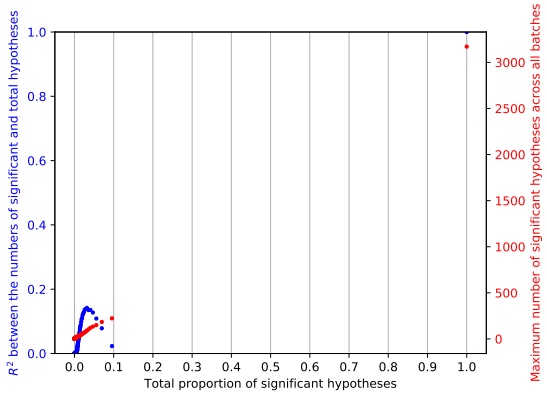
R^2 between the numbers of significant and total hypotheses
Maximum number of significant hypotheses across all batches
Total proportion of significant hypotheses

(c) `covTest` on null dataset

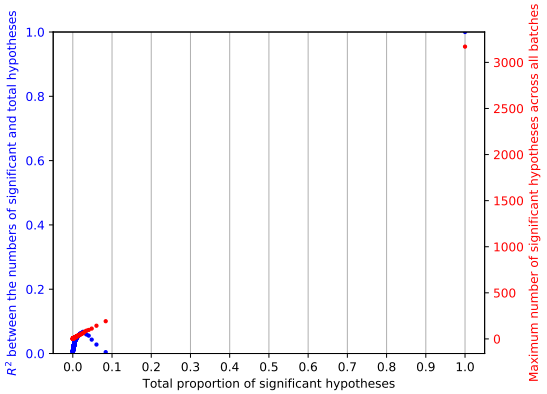
R^2 between the numbers of significant and total hypotheses
Maximum number of significant hypotheses across all batches
Total proportion of significant hypotheses

(d) `covTest` on real dataset

Figure 10: The R^2 goodness of fit is shown in blue as a function of the proportion of significant hypotheses for the high-dimensional null and real datasets. The maximum number of significant variables across all selection tasks is also drawn in red. Every dot corresponds to a significance threshold taken. Since the number of significant predictors can saturate in high-dimensional setting, Figure 10(a) and Figure 10(b) are only interpretable in region before the maximum number of significant variables plateaus.

In summary, `lassopv` again demonstrated optimal unified significance measure and FDR control and consequently more accurate p-values in all the tests on high-dimensional datasets, especially at small p-values which are of interest. Still, errors at large p-values were due to lasso itself and therefore unavoidable. Meanwhile, `covTest` failed to provide a uniform p-value distribution or unified significance measure, as illustrated by the four tests, whilst `selectiveInference` refused to deal with the high-dimensional datasets.

5 Discussion

In this article, we considered the practical problem of obtaining a unified significance measure for hypothesis tests in different batches. On one hand, we proposed statistical tests to evaluate the quality of p-value which were also generalized to any statistic as a unified significance measure. This would enable users of the p-value of any statistic to perform extensive evaluations on its quality with current or custom datasets. On the other hand, we proposed a lasso-based statistic for variable selection and computed its p-value. This would provide a discovery-oriented p-value measure for variable selection that allows downstream statistical treatments and is compatible with high-dimensional datasets. The R package, `lassopv`, which only requires a single computation of the lasso path and permits widespread adoption in practice, is available at https://github.com/lingfeiwang/lassopv.

Nevertheless, the accuracy of statistical inference is a different question from p-value quality. The statistical power of inference depends on many factors, e.g. the noise of data, the model or statistic applied for inference, and the quality of p-value if involved. Therefore, a higher quality of p-value in no means guarantees more accurate inference, and neither does `lassopv`. The inference power of `lassopv`'s statistic is beyond the scope of this paper and deserves future study.

Funding

This work has been supported by the BBSRC (grant numbers BB/J004235/1 and BB/M020053/1).

A Proofs

A.1 Proposition 1

First prove

$$\hat{\beta}_i^{(0)}(\lambda) > 0 \iff \lambda < \sigma_i^{(0)} \sigma_{\mathbf{y}_{\text{res}}}(\lambda) \text{Cor}(\mathbf{x}_i^{(0)}, \mathbf{y}_{\text{res}}(\lambda)). \tag{15}$$

Define

$$\tilde{\beta}_j^{(0)}(\lambda) \equiv \begin{cases} \alpha, & \text{for } j = i, \\ \hat{\beta}_j^{(0)}(\lambda), & \text{else.} \end{cases} \tag{16}$$

Since

$$\hat{\beta}_i^{(0)}(\lambda) \equiv \underset{\alpha}{\text{argmin}} \frac{1}{2n} ||\mathbf{y} - \mathbf{X}\tilde{\beta}^{(0)}||_2^2 + \lambda ||\tilde{\beta}^{(0)}||_1, \tag{17}$$

then

$$\hat{\beta}_i^{(0)}(\lambda) > 0 \tag{18}$$

$$\iff \frac{\partial}{\partial \alpha} \left(\frac{1}{2n} ||\mathbf{y} - \mathbf{X}\tilde{\beta}^{(0)}(\lambda)||_2^2 + \lambda ||\tilde{\beta}^{(0)}(\lambda)||_1 \right) \Big|_{\alpha=0^+} < 0 \tag{19}$$

$$\iff \frac{1}{n} \mathbf{x}_i^{(0)} \cdot (\alpha \mathbf{x}_i^{(0)} - \mathbf{y}_{\text{res}}(\lambda)) + \lambda \Big|_{\alpha=0^+} < 0 \tag{20}$$

$$\iff \lambda < \sigma_i^{(0)} \sigma_{\mathbf{y}_{\text{res}}}(\lambda) \text{Cor}(\mathbf{x}_i^{(0)}, \mathbf{y}_{\text{res}}(\lambda)). \tag{21}$$

Similarly,

$$\hat{\beta}_i^{(0)}(\lambda) < 0 \iff \lambda < -\sigma_i^{(0)} \sigma_{\mathbf{y}_{\text{res}}}(\lambda) \text{Cor}(\mathbf{x}_i^{(0)}, \mathbf{y}_{\text{res}}(\lambda)). \tag{22}$$

Combining the two, we get

$$\hat{\beta}_i^{(0)}(\lambda) \neq 0 \iff \lambda < \sigma_i^{(0)} \sigma_{\mathbf{y}_{\text{res}}}(\lambda) \left| \text{Cor}(\mathbf{x}_i^{(0)}, \mathbf{y}_{\text{res}}(\lambda)) \right|. \tag{23}$$

A.2 Proposition 2

Due to the rotational $SO(n)$ symmetry in the PDF of $\mathbf{x} \in \mathbb{R}^n$, each of which follows i.i.d $N(0,1)$, the distribution of $n\sigma_{\mathbf{x}}^2 \text{Cor}^2(\mathbf{x}, \mathbf{y})$ does not depend on $\mathbf{y}$. For simplicity, choose

$$y_1 = \sqrt{n/2}, \qquad y_2 = -\sqrt{n/2}, \qquad y_i = 0, \text{ for } i = 3, \ldots, n. \tag{24}$$

Then, by expanding the correlation, we have

$$n\sigma_{\mathbf{x}}^2 \text{Cor}^2(\mathbf{x}, \mathbf{y}) = (x_1/\sqrt{2} - x_2/\sqrt{2})^2. \tag{25}$$

Since $x_1, x_2 \sim i.i.d\ N(0,1)$, define

$$z \equiv x_1/\sqrt{2} - x_2/\sqrt{2} \sim N(0,1). \tag{26}$$

Therefore

$$n\sigma_{\mathbf{x}}^2 \text{Cor}^2(\mathbf{x}, \mathbf{y}) = z^2 \sim \chi^2(1). \tag{27}$$

References

[1] Carlo E Bonferroni. *Teoria statistica delle classi e calcolo delle probabilita*. Libreria internazionale Seeber, 1936. 2

[2] Olive Jean Dunn. Estimation of the medians for dependent variables. *The Annals of Mathematical Statistics*, 30(1):192–197, 1959. 2

[3] Olive Jean Dunn. Multiple comparisons among means. *Journal of the American Statistical Association*, 56(293):52–64, 1961. 2

[4] John D. Storey. A direct approach to false discovery rates. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(3):479–498, 2002. 2, 8

[5] John D. Storey and Robert Tibshirani. Statistical significance for genomewide studies. *Proceedings of the National Academy of Sciences*, 100(16):9440–9445, 2003. 2, 8

[6] Ronald Aylmer Fisher. *Statistical methods for research workers*. Genesis Publishing Pvt Ltd, 1925. 2

[7] R. A. Fisher. Questions and answers. *The American Statistician*, 2(5):30–31, 1948. 2

[8] Jonathan Taylor and Robert Tibshirani. Post-selection inference for l1-penalized likelihood models. *arXiv preprint arXiv:1602.07358*, 2016. 2, 3, 8

[9] Richard Lockhart, Jonathan Taylor, Ryan J. Tibshirani, and Robert Tibshirani. A significance test for the lasso. *Ann. Statist.*, 42(2):413–468, 04 2014. 2

[10] Jason D. Lee, Dennis L. Sun, Yuekai Sun, and Jonathan E. Taylor. Exact post-selection inference, with application to the lasso. *Ann. Statist.*, 44(3):907–927, 06 2016. 2

[11] Ali Shojaie and George Michailidis. Penalized likelihood methods for estimation of sparse high-dimensional directed acyclic graphs. *Biometrika*, 97(3):519–538, 2010. 2, 7

[12] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996. 2

[13] Rina Foygel Barber and Emmanuel J. Candès. Controlling the false discovery rate via knockoffs. *Ann. Statist.*, 43(5):2055–2085, 10 2015. 3

[14] Tuuli Lappalainen, Michael Sammeth, Marc R. Friedlander, Peter A. C. /'t Hoen, Jean Monlong, Manuel A. Rivas, Mar Gonzalez-Porta, Natalja Kurbatova, Thasso Griebel, Pedro G. Ferreira, Matthias Barann, Thomas Wieland, Liliana Greger, Maarten van Iterson, Jonas Almlof, Paolo Ribeca, Irina Pulyakhina, Daniela Esser, Thomas Giger, Andrew Tikhonov, Marc Sultan, Gabrielle Bertier, Daniel G. MacArthur, Monkol Lek, Esther Lizano, Henk P. J. Buermans, Ismael Padioleau, Thomas Schwarzmayr, Olof Karlberg, Halit Ongen, Helena Kilpinen, Sergi Beltran, Marta Gut, Katja Kahlem, Vyacheslav Amstislavskiy, Oliver Stegle, Matti Pirinen, Stephen B. Montgomery, Peter Donnelly, Mark I. McCarthy, Paul Flicek, Tim M. Strom, The Geuvadis Consortium, Hans Lehrach, Stefan Schreiber, Ralf Sudbrak, Angel Carracedo, Stylianos E. Antonarakis, Robert Hasler, Ann-Christine Syvanen, Gert-Jan van Ommen, Alvis Brazma, Thomas Meitinger, Philip Rosenstiel, Roderic Guigo, Ivo G. Gut, Xavier Estivill, and Emmanouil T. Dermitzakis. Transcriptome and genome sequencing uncovers functional variation in humans. *Nature*, 501(7468):506–511, 09 2013. 3, 7

[15] Ryan J. Tibshirani. The lasso problem and uniqueness. *Electron. J. Statist.*, 7:1456–1490, 2013. 5

[16] Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. Least angle regression. *Ann. Statist.*, 32(2):407–499, 04 2004. 5, 12

[17] Jerome Friedman, Trevor Hastie, and Rob Tibshirani. Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, 33(1):1–22, 2010. 5

[18] Nir Friedman. Inferring cellular networks using probabilistic graphical models. *Science*, 303(5659):799–805, 2004. 7

[19] Bryon Aragam and Qing Zhou. Concave penalized estimation of sparse Bayesian networks. *arXiv preprint arXiv:1401.0852*, 2014. 7

[20] Eric E Schadt, John Lamb, Xia Yang, Jun Zhu, Steve Edwards, Debraj GuhaThakurta, Solveig K Sieberts, Stephanie Monks, Marc Reitman, Chunsheng Zhang, et al. An integrative genomics approach to infer causal associations between gene expression and disease. *Nature Genetics*, 37(7):710–717, 2005. 7

[21] Lingfei Wang and Tom Michoel. Efficient causal inference with hidden confounders from genome-transcriptome variation data. *arXiv preprint arXiv:1611.01114*, 2016. 7